

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
13 March 2003 (13.03.2003)

PCT

(10) International Publication Number
WO 03/020112 A2

(51) International Patent Classification⁷: **A61B**
(21) International Application Number: **PCT/US02/27586**
(22) International Filing Date: **30 August 2002 (30.08.2002)**
(25) Filing Language: **English**
(26) Publication Language: **English**
(30) Priority Data:
60/315,953 **30 August 2001 (30.08.2001)** **US**

(71) Applicant (for all designated States except US):
PHILADELPHIA OPHTHALMIC IMAGING SYSTEMS [US/US]; 1208 Henry Lane, Gladwyne, PA 19035 (US).

(72) Inventors; and
(75) Inventors/Applicants (for US only): **SINCLAIR, Stephen, H. [US/US]; 1208 Henry Lane, Gladwyne, PA 19035 (US). BHASIN, Sanjay [US/US]; 953 Talealamore Drive, Ambler, PA 19002 (US).**

(74) Agents: **BARKER, Bruce, J. et al.; Pennie & Edmonds LLP, 1155 Avenue of the Americas, New York, NY 10036 (US).**

(81) Designated States (national): **AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZM, ZW.**

(84) Designated States (regional): **ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, SK, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).**


Published:

— without international search report and to be republished upon receipt of that report

[Continued on next page]

(54) Title: **SYSTEM AND METHOD FOR SCREENING PATIENTS FOR DIABETIC RETINOPATHY**

Threshold Level: **72**



Threshold Level: **55**



Histograms of two images of the fundus.

(57) Abstract: The present invention comprises a robust technique to automatically grade retinal images through detection of lesions that occur early in the course of diabetic retinopathy: dot hemorrhages or microaneurysms, blot and striate hemorrhages, lipid exudates and nerve-fiber layer infarcts. In addition, the present invention includes methods to extract the optic nerve in the appropriately identified fields, and to track and identify the vessels (measuring vessel diameters, tortuosity and branching angles). The present invention preferably identifies 3 levels: no retinopathy, microaneurysms alone, and lesions additional to microaneurysms, the latter two levels being the earliest detectable forms of the disease. The method of the present invention is particularly suited to overcoming the difficulties in grading retinopathy, which stems from image-to-image variation, low contrast of some the lesions from the background and non-uniform lighting and flare within the same image resulting in variation in the different quadrants of the same image. The system has the capability of utilizing human expert marking of images and grading of retinopathy to improve image quality assessment, lesion detection and retinopathy grading.

WO 03/020112 A2

BEST AVAILABLE COPY



For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

System and Method for Screening Patients for Diabetic Retinopathy

Summary of the Invention

5

One of the complications of diabetes is the gradual degradation of vision known as diabetic retinopathy. It is known that over time, a majority of diabetics will lose some vision to this condition, and the present state of the medical art is to treat the retinal lesions that mark the disease with laser light. Diagnosing diabetic retinopathy, however, requires a trained specialist to view the retina (or a photograph of the retina) at periodic examinations, and to recognize small lesions and changes therein. Because there are many more diabetics than specialists, and because they must be examined at regular intervals, difficulties and deficiencies have arisen in health-care systems that attempt to screen and treat diabetic retinopathy. The present invention provides an efficient computer-implemented screening analysis of retinal photographs to identify the stages of diabetic retinopathy.

10

15

Background of the Invention

Diagnosis of diabetic retinopathy is commonly performed by skilled medical personnel (usually ophthalmologists) in direct examination of the fundus, or by evaluation of sets of fundus photographs taken using special-purpose cameras. Such examinations, whether done directly or by review of photographs is a time-consuming and inexact process, and experts in diagnosis often disagree in their results, particularly in the very earliest stages of the disease. In addition such diagnosis is expensive, and is often performed at intervals that are longer than desired due to the cost and lack of available skilled manpower.

20

25

Brief Description of the Invention

30

The present invention comprises a robust technique to automatically grade retinal images through detection of lesions that occur early in the course of diabetic retinopathy: dot hemorrhages or microaneurysms, blot and striate hemorrhages, and

lipid exudates. In addition, the present invention includes methods to detect nerve fiber layer infarcts, extract the optic nerve in the appropriately identified fields, and to track and identify the vessels (measuring vessel lumen diameters, tortuosity, and branching angles). The present invention preferably identifies 3 levels: no
5 retinopathy, microaneurysms alone, and lesions additional to microaneurysms, the latter two levels being the earliest detectable forms of the disease. The method, however, may be expanded to 7 levels through the detection of the lesions that occur in advanced stages of the disease and may be utilized, through the evaluation of changes of the normal vasculature (lumen diameter, tortuosity, and branching angle),
10 to evaluate the risk of developing retinopathy.

The method of the present invention is particularly suited to overcoming the difficulties in grading retinopathy, which stems from image-to-image variation, low contrast of some the lesions from the background and non-uniform lighting and flare
15 within the same image resulting in variation in the different quadrants of the same image.

The expert decision system implemented in the present invention bases the retinopathy grade determination upon the results of the lower-level detectors (lesion
20 detection) for each eye photographed. The system may be tuned for individual detectors and for mydriatic images and non- mydriatic images using separate parameters based on the camera type, image type, characteristics of patients, and characteristics of fundus.

25 A data archive is the core of the data management system that acts as the central repository for all patient data: images, demographics as well as reports. This archive will be accessible for storage and retrieval via the Internet and will be of a scalable design to accommodate growth over time. The benefits of a centralized data management architecture include (1) The ophthalmologist or retina specialist will
30 have access to current as well as prior studies for comparison regardless of where the historical data was acquired. (2) A central data repository will allow for an objective and quantitative means to evaluate the progression of the disease in individuals or populations over time in terms of the changes of normal vessels, the change of

existing diabetic lesions ,and the occurrence of new ones. This database will provide the foundation to develop regression and serial studies to develop risk prediction algorithms in the future. (3) Algorithms that will scan the archive will produce quantitative measures of vascular tortuosity, branching angle, and caliber variation, which have been identified as markers of vascular disease and which can be tracked over time. Again, this will also enhance risk prediction, predominantly in the early stages, prior to or subsequent to the development of retinopathy. It should be noted that these parameters are not able to be assessed by human grading. (4) Data mining of the massive warehouse of data can allow screening proficiency and patient compliance to be examined, as well as providing valuable insight into the trends in the various populations and a comparison to observe and follow over time the effectiveness of treatment among patient populations (e.g. similar patients in each clinic).

Brief Description of the Figures

Figure 1 depicts histograms of two images of the fundus.

Figure 2 depicts an initial fundus image, and the result of retina extraction processing.

Figure 3 depicts the result of processing the retina image by coalescent filter [Bhasin and Meystel 94] to obtain the vessels.

Figure 4 depicts graphically the problem with locating the retina in a photograph.

Figure 5 depicts inseting.

Figure 6 (a) depicts a subsampled original retinal image.

Figure 6 (b) depicts the image of Figure 6(a) after subtraction of vessels from the image.

Figure 7 depicts results of finding the retina using geometry to improve the pickup.

Figure 8 depicts an overall flow diagram for the method of the present invention.

Figure 9 depicts a flow diagram of the image processing method of the present invention.

Figure 10 depicts a flow diagram of the photographer assistance method of the present invention.

Figure 11 depicts a flow diagram of the serial (over time) study method of the present invention.

Detailed Description of the Invention

According to the present invention, images are acquired locally with immediate
5 feedback to photographer to assist in optimizing each image quality using a small
looping process to guide the photographer, analyzed locally and stored locally on CD-
ROM and magnetic storage. Screening is reported immediately from the system with
a printed report. Those eyes that "fall out" of screening, either because a significant
10 number of photographs cannot be analyzed (at present too poor despite 3 attempts to
photograph) for one eye or because they are above the accepted threshold in
number/type of lesions/level of retinopathy, will be transmitted to a remote specialist
ophthalmologist for reviewing. The specialist will then transmit back a review with
grading and recommendations: either the patient is to be advised to make an
15 appointment for further examination, or the photographs are sufficient but all the eye
needs is repeat screening at a recommended interval. This report is maintained in the
database for data mining with regard to the success and thresholds for screening and
referral.

20 In practicing the present invention, images may be obtained in one of the following
ways:

- (a) Monochromatic 30° images digitized from 35 mm Ektachrome® slides at at
least 1024 X 1024 X 8-bit depth using a green Wratten filter
- (b) Photograph eyes with a digital on-line system using a fundus camera with a
digital camera back allowing color photography at at least 2 meg x 2 megx 32 bit
25 depth resolution
- (c) Photograph eyes with a digital on-line system using a fundus camera with one
or more interposed filters including 535nm notch filter, 605 nm notch filter (or
others), a digital camera back allowing monochromatic 1024 X 1024 X 8-bit depth

30 The present invention preferably uses a protocol of 5 photographs of 35° fields per
eye.

The present invention also provides for cataloging patients (including naming images as well as providing other demographic and clinical data) to enable batch processing and historical tracking of disease as well as identification of risk indicators. It has been determined that cataloging patients is important: keeping digital photos on file is efficient for doctor's office – for example, images may be maintained in an image-management database which also allows the recording of ancillary patient data and physician notes for each set of photographs.

The image-management system of the present invention allows medical professionals to track improvement and determine more accurately what factors in what patient classes dictate degradation rate, which dictates patient recall and reexamination frequency.

The system of the present invention uses a file-naming convention of the images that allows the system to automatically locate images of the same patient and the corresponding field. The naming scheme for all images is: PatientID- Camera type.Eye-Field-Field of view.?Image type.Processing

- Eye is either L or R
- Field is a number from 1 thru 5
- PatientID is system generated alphanumeric identifier
- Camera type is a number 1 for non-mydratic camera, 2 for mydratic camera
- Field of view is the angular degree of retina photographed by camera with each image (i.e. 30° or 35° or 45°)
- Image type is a number for the type of image: 1 for digital direct color image, 2 for monochromatic derived from digitization from Ektachrome slides, 3 for monochromatic followed by the notch filter peak: (e.g. 535 for a 535nm peak filter).

Processing is a code that indicates the type of the image: RAW (for original photograph), GRA (graded image indicating lesions present).

The method of the present invention includes a software system for assisting the photographer in capturing adequate retinal images. In use, the photographer photographs the eye and indicates the field being photographed. When these images are captured, each photograph is marked with an identifier: the current system utilized

but not to be restricted i.e. OD field#1 centered on fovea with the disc on the right or left of the field depending up on the eye; OD field #2: disc in lower left (superonasal) OD field #3: disc in upper left (inferonasal); OD field #4 disc in upper right corner, fovea in upper center (inferotemporal; and OD field #5 disc in lower right corner, fovea in lower center (superotemporal); the OS fields are just the opposite, right from left. These 5 photographs will give some overlap between the 45 degree images with several photographs taken of the fovea and span all of the necessary areas of the retina that require photography and examination. It is known that the present invention may be improved if the photographer marks the fovea in field #1. Each photograph is then:

- Examined for these necessary elements, disc and vessels and their position in the photograph etc.
- Contrast: The background is compared with the vessels (loss of contrast is produced by moving the camera too far in or too far out- on the Canon CR6, for example, this is helped by having the photographer align some bright alignment dots.
- Focus: de-focused image is detected by examining edge definition of central vessels.
- Alignment (Up/down/right/left within the pupil): Misalignment with regard to the pupil produces an edge flare within the field of the photograph (loss of contrast):- this will require contrast evaluation at the edge of the field as well as in the center to detect the flare (to be compared with a database of some example images, taken but yet to be determined) and tell the photographer to realign if the edge flare is detected.

If the photograph is determined to be inadequate for any of the parameters examined, a set of instructions for each element is immediately presented to the photographer with instructions about which field and eye to re-photograph, and how to improve the image. This leads, not only to improved quality of images acquired for an individual but also leads in general to more rapid improvements in retinal photography by the inexperienced photographer.

Retinopathy Grading

Retinopathy grading in the present invention is based on image processing and computer vision software for lesion detection and quantification from digitized monochromatic images taken either on line with 1280X1024 pixel density digital monochromatic camera, or utilizing the green channel from a trichromatic color camera with similar one channel pixel density, or digitized at similar resolution from slide film images at 30⁰ or 50⁰ taken from a non-mydriatic camera. Reference is made to the CD-ROM Appendix which contains the computer source code of the preferred embodiment of the present invention.

Quantification: (for each field) number /field, total area/ field area evaluated, histogram of number vs size, histogram of number vs. density in each field.

Screening results in 3 levels: 1) no lesions (hemorrhages, lipid exudates, or nerve fiber layer infarcts), 2) microaneurysms only (dot hemorrhages), 3) dot hemorrhages with other hemorrhages, exudates or infarcts.

The separate detectors are for the following lesions and anatomical features in the retina:

- A. Hemorrhages: dot, blot, striate
- B. Lipid Exudates
- C. Nerve fiber layer infarcts
- D. Optic Nerve Head
- E. Vessels (Arteries, Veins): Detection of 1⁰, 2⁰, 3⁰ branching order

Other lesions of more advanced retinopathy grades for which detectors are being pursued:

- A. Intra-retinal microvascular anomalies (IRMA)
- B. Venous loops
- C. Epi-papillary neovascularization
- D. Epi-retinal neovascularization
- E. Sub-hyaloid or vitreal hemorrhage
- F. Epi-retinal fibrosis

G. Retinal detachment

Approaches to determine grade include:

1. Geometric techniques
- 5 2. Image Processing to obtain Position, Size of Lesions
3. Matched Filter - Find something that is oval - use Biology code
4. Morphological Filter - Grow-Shrink method
5. Remove vessels -
 - Find all junctions
 - 10 - Find vessels emanating from them and follow them

Region of Interest Extraction

For images digitized from film, the portion containing the retina is extracted. For both digital and digitized images, the bad contrast portions caused due to flare, etc.
15 are removed. An improvement to the present invention removes certain portions of photographs which result from image flare and contain no useful data.

The histograms of two different fundus images shown in Figure 1 demonstrate that there is a considerable separation in gray values between the background and the
20 retina. Unfortunately this separation occurs in a broad range of values. There are other factors such as flare at the edges of the retina and artifacts in the imaging that severely limit the use of a simple technique.

These vessel boundaries will be used to determine the primary, secondary and tertiary
25 vessels. An adaptive threshold is used to locate the vessel. Its settings depends on image contrast. This requires locating the retina within the image and measuring its contrast. To discriminate the retina from the background the inventors have used the ISODATA clustering technique to first define a threshold value and then binarized at that value. The ISODATA algorithm is an iterative method based upon the gray level
30 histogram of the image. The histogram is split up into two parts, the foreground pixels and the background pixels, assuming an initial threshold value. Then the average value of the foreground and of the background pixels is calculated and a new threshold value is taken midway between these two average values. This process is

repeated, based upon the new threshold estimate until the threshold value does not change any more.

It often happens that this does not result in definition of the full retina because of the presence of a deep shadow over part of the image (Figure 4). One way to overcome this obstacle is to use a restricted convex hull. First the largest object is found and then all the combinations of two contour points with a Euclidean distance less than or equal to some preset distance (d), are connected by a straight line. If a background pixel is found on such a line, it is added to the original object. This operation also closes all holes which are less than d pixels wide.

An alternative approach is not to use the convex hull, which creates the need to close off any internal holes within the segmented retina. To do so, generate a binary image using the ISODATA technique and remove all but the largest object. Then invert the binary image and remove all objects that are not contiguous with the boundary of the image, that is that are not part of the background. Then re-invert the binary image.

Once the retina is located, inset the retina mask edge to define an inner boundary (Fig 5). An adequate inset distance for the 1k images is 100 pixels. The reason for performing this operation is to avoid including the retinal edge in measuring the retina contrast. Often the edge is subject to deep shadow or bright flare.

The histogram of the retina mask following insetting is used to define the region of the image for which the gray value histogram is determined. From the histogram the standard deviation is computed and that is used as a measure of the image contrast. The rank order to be utilized in the obtaining of the vessels needed to be derived. Different ranks were used and the results are shown for rank = 0.25, 0.5 (median filter) and 0.75 in Figure 6. Here the original image was a subsampled version of the scanned image to reduce the processing time. The subsampling was by a factor of 4. The subsampling was done by a straightforward linear method and other techniques like bilinear sampling might improve the results and will be studied later. To see the effectiveness of the technique, the vessels found were subtracted from the original image (subsampled scanned image by a factor of 4). The original image and the

results of the subtraction are presented in Figure 7. We see that there are a lot of false alarms which can be isolated by form factor and total area of object.

Discrimination of the vascular bed from the background is accomplished on the basis of density, size and shape. The algorithm utilizes an adaptive threshold based on a rank order filter. The routine assigns a value of 0 (background) or 1 (object) to a given pixel depending on its value relative to that of its neighborhood: All the pixels of the neighborhood are rank ordered on the basis of gray value, a constant offset is added to the value of the central pixel and the result compared to the value of the i^{th} element of the neighborhood. If it is lower, that is the central pixel is substantially darker than its surround, then the central pixel is assigned a value of 1, else it is given a value of 0. Three interrelated parameters important to the operation of this adaptive threshold are the offset, kernel size and count threshold (i). The kernel size and offset are constants both for a given image and across images. We have tested both circular and rectangular kernels. As you would predict, the circular ones result in less artifact. We settled on a circular filter with a diameter of 17 pixels. The offset used was 15. The count threshold is kept constant for a given image but is varied from image to image depending on the image contrast.

To set the count threshold the standard deviation of the retinal image is used. To assess how the count threshold varies with the image contrast, the standard deviation of the histogram of each of twenty-five images was computed and each image was convolved with the adaptive threshold with the count threshold varied until the best segmentation was achieved as judged visually. The relation between the two has been modeled by a linear fit ($y=234.34-1.607x$; $r = 0.88$; Fig. 3).

Vessels identified and removed from image to leave only dot hemorrhages as dark structures. The present invention can remove objects like the optic nerve head which are very bright to get a clearer representation of the gray value distribution (histogram) in the rest of the image. Much of the processing is dependent on grey values and these techniques along with specialized algorithms like the k-means, modified gradient, signature analysis and morphology allow us to perform adaptive segmentation that leads to a more reliable extraction of the lesions.

The present invention applies a smoothing filter – special configuration that removes noise but does not destroy the structure of the hemorrhages or cause artifacts in the image.

5

The lesions detected by the various detectors are fed into an expert system of known construction. The expert system is constructed out of rules built from consultations with expert ophthalmologists. It embodies the reasoning used by them to make interpretations. Among the features queried of the experts and rule-coded are:

10

1. Quick scan of the image that detects of vessel fragments.
2. Morphological processing to connect co-linear fragments and branches to the major vessel (dilate and then constrict the regions)
3. Skeletonization algorithm developed for finding the median line
- 15 4. Subtract the vessels found from the original image to see efficacy of processing

20

The present invention uses a path-searching technique. To find the path, we start at the position in the gray value image corresponding to a given endpoint and search for the direction in which the mean gray value is maximum. The initial search direction is determined by examining the skeleton and the connectivity of the endpoint and setting the initial direction (d) to the opposite direction. Given a current pixel (x, y) and a current direction (d), there are three candidate points for the next pixel, viz. the neighbor points in the directions (d-1), (d) and (d+1) module 8. To decide which to take, three straight lines are considered, one in the direction (d), another in between the directions (d) and (d+1) and a third in between the directions (d) and (d-1). Along each of these lines the mean value is calculated, over a length of <val> pixels. The next pixel is the candidate point that corresponds with the direction with the maximum average. The process is repeated at the new pixel, using the found direction as a new value for (d). The process stops, firstly, if the distance of the candidate pixel to the image boundaries is less than a preset value. Further, if the maximum of the mean values in the candidate directions is less than the specified value <minval>. Next, if a pixel is obtained that is set in the output bitplane. For this the system uses

25

30

the background outside the eye. Finally the process is stopped if the number of pixels found becomes equal to a preset value.

Alternatively examination of the gradient normal to the generated curve at a regular
5 spacing may be used to reject the curve.

Gradient

```
CStatistics stats(m_doc, original);
CBasicFilter bf(m_doc, m_src); bf.Sobel();
10 CStatistics stats2(m_doc, m_src);
BYTE thr = (BYTE)( stats2.m_mu + k1* stats2.m_sd);
bf.Threshold( thr );
Copy either m_src->m_data[i] or 0 to m_dest depending on
contrast measure of background (to eliminate bright flare)
15 if( ref->m_data[i] > (stats.m_mu + k2* stats.m_sd) )
    m_dest->m_data[i] = 0;
else
    m_dest->m_data[i] = m_src->m_data[i];
```

20 Methodology for Detectors

Separate the veins from arteries, measure vessel branching patterns, vessel tortuosity, vessel caliber variation. Locate the disc and fovea so that lesions can be localized in position relative to the fovea.

25 General philosophy: Encode tests as functions that return three possible scores:

1 (SUCCESS), 0 (Indeterminate or DONT_KNOW) or -1 (FAIL)

This setup will be used later to multiply output of functions with weights and combine to produce a confidence score.

30 As discussed above, following the adaptive threshold and noise removal, vessel segmentation is achieved on basis of object size and form factor. The objects rejected

are in fact classified as hemorrhages. Size is used to distinguish dot from blot hemorrhages (DH & BH, respectively) with later being the larger of the two.

Adaptive Segmentation

- 5 Run K-means on large image segments - 4x4 or 8x8 segments in whole image - to produce three regions: background, dark regions like vessels and dot hemes, and light regions like lipid exudates. This will give a rough segmentation of the image that takes the local shading and variation into account. The mean grey values of the regions will be used later as pivotal values for finer segmentation.

10

Technique for hemes

If any 2 of following tests are failed, i.e. if (test1() + test2() + ... == -2)

0. Rank order the grey values in a NxN region and then divide them into M-tiles (if M = 5, it would be quintiles). The pixels in the lowest rank order are colored RED.

15

- For a dot heme, the RED part should be somewhat in the center as the ideal profile of the heme is a crater – low grey values in the center surrounded by concentric rings of larger value.
- A lipid exudate is like a mound which has the same characteristics as heme for an inverted image.

20

1. Centre 4x4 of full stamp does not have any RED
2. Let XMIN, YMIN, XMAX, YMAX be corners of full stamp and (ymin, xmin1, xmax1) be first ival of RED, and (ymax, xmin2, xmax2) be last ival -- these form actual corners of RED. Check if any edges of RED are at corners of full as in the case for triangles formed in peripheral segments of image and in flare regions As you have a regionarray, XMIN = lesionpost->supp.x1, ... so test is

25

(lesionpost->supp.x1 - xmin1 > -2 and lesionpost->supp.y1 - ymin > -2) OR
 (lesionpost->supp.x1 - xmin2 > -2 and lesionpost->supp.y2 - ymax < 2) OR
 (lesionpost->supp.x2 - xmax1 < 2 and lesionpost->supp.y1 - ymin > -2) OR
 (lesionpost->supp.x2 - xmax2 < 2 and lesionpost->supp.y2 - ymax < 2)

30

3. To check if it is a vessel, see if it runs for full length of stamp
 lesionpost->supp.y1 - ymin > -2 and lesionpost->supp.y2 - ymax < 2
4. If total size of stamp is > 60x60 return FAIL

5. Form factor of RED should be between 0.9 and 1.1 which eliminates trapezioids for vessels and triangles and allows circular objects

Technique for Cotton Wool Infarcts and Lipid Exudates

- 5 The CWS and EX are differentiated from the background of the fundus images by first inverting the image gray value so that the exudates are now darker than the background and then applying the same adaptive filter used to segment the hemorrhages. After removal of noise with a median filter the remaining objects are separated into CWS and EX. This is done on the basis of size and gradient. Objects
10 smaller than a fixed constant and with a sharper gradient are classified as EX. Remaining objects are classified as CWS. Obviously, the gradient is assessed on the original gray value image.

Refinements

- 15 Developed an approach to detect objects which are touching the vessels. It employs a distance transformation which replaces each pixel of an object with an estimate of its shortest distance to the background (the distance to the nearest background pixel). To speed production of the map, the inventors have used pseudo-Euclidean distances:
20 distance to the nearest neighbor is 5, to the diagonal it is 7, and knight move is 11. Now thresholding at some value, for instance, 50, will capture pixels which are 10 pixels or further from the nearest edge of the object. In this way it is possible to detect an unusually thick vessel.
- A refinement of this approach looks not at the absolute thickness but rather the change
25 in thickness. Here we are looking for a sudden change from a thin to a thick vessel and then back down to a thin vessel. To do this the skeleton of the binary vessel image is found using a distance transform. It differs from the medial-axis transform in that here the skeleton is defined as a set of connected, one pixel thick arcs, lying midway between the object boundaries and being a topological retraction with the
30 same connectedness as the original object. Thus unlike the medial axis connectivity is preserved. It differs from the Hilditch skeleton in that the way it is applied, the distance skeleton is computed using pseudo-Euclidean distances rather than city-block distances (1 nearest neighbor, 2 diagonal, 3 knight move). Now, the skeleton is used

as a mask with the distance map to obtain distances from the skeleton to the edge of the object.

Level 1

5 **Filter and Run CLesion (Dot,Blot etc) with**
 IndicationFilter<DarkonLight/LightonDark>
 (Converts grayscale to binary with small blobs in image space).

10 Filter to smooth noise and generates vessel bed. The use of filtering is configurable and is determined through the use of experiments done on images from a specific camera that measure the relative merit.

15 Runs a filter with 3x3 cross at center and 8-neighbor 'arms' ending in 2 pixel long 'hands' at end of each arm. The center grey values (weighted sum) is compared with sum of 'hands'. If darker than all, hit is marked at center 3x3. Hands are moved out by StepSize from center, starting at minSize and ending at. Multiple passes were added to pick up more than one lesion size; maxSize. Is crude matched filter - many false alarms. Call with minSize & maxSize ODD, stepSize EVEN. Looks like:

20 * * *
 * * *
 @
 ** @@@ **
 @
 * * *
 25 * * *

Last step is to PruneBySize...

30 Converts CRA to object list and applies simple geometrical rules to prune objects in list. List is held and managed in CIArray - see definitions there. min and max threshold are used to specify color-range to consider for object e.g. red and green is from 1 to 2. Allows pruning by: # pixels constituting object, and min & max extent of object. Permits specification of AND/OR for x & y extent-limit rule and whether or not to FILL objects that are passed.

Level 2

Processes in object space. Initiates call to <CObjectsManipulation::PruneBySize>. PruneBySize is a common function for all lesions, has no separate body

5

BOOL CLesion::GeometricFilter(BOOL bUseDlg)

Converts CRa to a cohesive group of pixels or a list object and apply simple geometrical rules to prune the objects in the image by: # pixels constituting object, and min & max extent of object.

10

Compute first and second Moments(area, centroid, major and minor axis of best-fit ellipse)

Test1: area > minPix

15 Test2: Compare both xExtent & yExtent to minimum and max thresholds or check if only either of them meet the minimum constraints. Permits specification of AND/OR on the constraint i.e. either dimension or both dimension to be constrained for x & y extent-limit rule.

20

Level 3

Refines image-space and object-space evidence with help of expert rules.

CLesion::SignatureFilter(BOOL bUseDlg)

Modified 02/28/97 to allow processing based on cf5 OR raw signature. Major overhaul of CDialogManager also done, and ini file also. 3 new stringsections per lesion ADDED to ini and 1 DELETED i.e. L3[lesion]Dlg replaced by L3[lesion]CommonDlg, L3[lesion]Cf5Dlg, and L3[lesion]RawDlg. First Dlg determines whether user wants to use cf5 or raw (BOOL, m_param10) and other parameters common to both methods. Second dialog is selected on the basis of value of m_param10 and queries for those parameters which are specific to the particular method selected earlier.

25
30

CLesion::SignatureFilterBasedOnRaw

Function SPLIT 2/27/97 to allow different processing of two types - original (using cf5) or new (using raw) for analysis Function modified 1/17/97 to create diagnostic file with fixed name in current results directory. This file contains yellow filled, passed objects for final analysis. File is overwritten on each pass through. Valid
 5 during batch processing of a single image.

PROCESS

```

int reason;
BOOL bPassed;
10 CScreenerApp* pApp = (CScreenerApp*)AfxGetApp();
   CRegionarray* refCRa = new CRegionarray(m_doc, m_doc->m_pFiles.raw)
   CRegionarray* vslCRa = new CRegionarray(m_doc, m_doc->m_pFiles.vsl)
   CRegionarray* diagCRa = new CRegionarray(m_doc, refCRa->m_dX, refCRa-
   >m_dY)
15   CIList ci(m_doc)
   CIArray ca(m_doc, m_doc->m_pImage, minThr, maxThr)Get objects formerly in
   "biglist"

20   while( ca.GetNext( &ci ) )
   {
   BOOL bNotOnVessel = NewNotOnVessel(&ci, vslCRa, &bPassed, &reason); 1 Pixel
   overlap
   if( bNotOnVessel || (m_lesionType==IBLOT) )
25   {
   CExpertSystem ce(m_doc, m_lesionType, bNotOnVessel, &ci, refCRa, userbins,
   borderWidth, bUsePercentile, nObjectsMax, fillFactorMin, formFactorMin,
   fillFormMin, pcntAreaMin, tr11, tr21, tr31, tr12, tr22, tr32, bUseLocalContrast,
   minGValLT, maxGValGT, minDRange, maxDRange, minBinWidth1,
30   minBinWidth2, minBinWidth3, minBinWidth4);

   if( (bPassed = ce.Verify(&reason)) == TRUE )
   {

```

```
InsertLesionInOutputImage(&ci, diagCRa, ce.m_analysis->m_stamp,
bStampBorderOnly, bStampOneColor);
```

```
ce.UpdateDocStatistics();For each incremental PASSED lesion
```

```
5 record info for entry in database
}
```

```
WriteOutputImage(diagCRa);
```

```
10
```

```
15 Default constructor creates an array of SIGNATURESTATISTICS that is currently
HARDCODED to size = 5; Of these index 0 is a placeholder for statistics of the
stamp's CRegionarray parameters and 1-4 holds statistics for the R/G/Y/B stamps
(islands in regionarray). In verification, it will be necessary to guarantee that statistics
of each color have meaning. NULL structs in the array will be checked there.
```

```
20 CExpertSystem::CExpertSystem(CScreenerDoc* pDoc, int lesionType, BOOL
notonvsl,
CIList* ci, CRegionarray* refCRa, int userBins, int borderWidth, BOOL
bUsePercentile, int nObjectsMax, double fillFactorMin, double formFactorMin,
double fillFormMin, double pcntAreaMin, int tr11, int tr21, int tr31, int tr12,
int tr22, int tr32, BOOL bUseLocalContrast, int minGValLT, int maxGValGT,
25 int minDRange, int maxDRange, int minBinWidth1, int minBinWidth2, int
minBinWidth3, int minBinWidth4)
```

```
30 1. Instantiate CStampAanalysis (which looks at <RED if dot/blot, YELLOW
lipid/cwool (OIS), BLUE lipid/cwool (Kaiser)> part of signature)
```

```
2. Initialize stats
```

```
typedef struct _SIGNATURESTATS
{
```

shortnObjects;# of objects (of a particular color) in a given stamp (CIArray)

doubleaveSize;average size

doublesizeSD;modal size

doubledispersion;centroid of areas normalized by center x,y

5 doublelargestArea;pixels in largest object

CRectlargestRect;rect of largest object

}SIGNATURESTATS;

SIGNATURESTATS m_ss.SetSize(5);Array of SignatureStatistics structures

10

Initialize:m_ss[0].nObjects, .largestArea, .aveSize, .sizeSD, .dispersion,

.largestRect.left, .largestRect.right, .largestRect.top, .largestRect.bottom

(This corresponds with stamp Rect, the other numbers are not used)

15 GetSignatureStatistics

Function added 030497 to assist in designing an expert system for signature analysis.

SIGNATURESTATS definition is in CScreenerApp. Parameters to be added till

accuracy of the damned system is adequate - or unimprovable. May split into multiple

hierarchical functions if too much redundant processing becomes necessary with a

20 single function. Will keep a single SIGNATURESTATS structure in any case.

m_ss[1] = ((CIArray*)(m_analysis->m_CIArrays[0]))->GetSignatureStatistics(

((m_analysis->m_stamp->m_x1+m_analysis->m_stamp->m_x2)/2),

(m_analysis->m_stamp->m_y1+m_analysis->m_stamp->m_y2)/2); DO RED

25

3. Do Signature Analysis Including

Investigates 8-neighborhood of core-color object to count and record the transitions to

next-color bands. In other words looks for deviation from ideal case which is

30

44444444

43333334

43222234

43211234

43211234

43222234

43333334

5 44444444

This cannot be done with list objects because "next-to" and "nearly-surrounds" are hard to compute. This function provides a working metric on a scale of zero to 8 for transitions from each band into the next, e.g. (7,6,3) which allows use in rules.

- 10 Function requires start point within core-color region of regionarray. Yet to determine effect of concave objects.

FAIL unconditionally if local contrast not high enough and sanjay wants to use this rule Added 070998

- 15 if(m_bUseLocalContrast) rule to apply test or skip it

```
{
```

```
if( m_stamp->m_bw1 <= m_minBinWidth1 )
```

```
{m_bSuccess = FALSE; m_reason = 22; return;}
```

- 20 if((m_stamp->m_gvMax - m_stamp->m_gvMin + 1) < m_minDRange)

```
{m_bSuccess = FALSE; m_reason = 23; return;}
```

```
if( (m_stamp->m_gvMax - m_stamp->m_gvMin + 1) > m_maxDRange )
```

```
{m_bSuccess = FALSE; m_reason = 23; return;}
```

- 25 }

FAIL unconditionally if largest object bleeds

```
m_bOnEdge = CheckEdge(m_ss[1].largestRect, m_ss[0].largestRect);
```

```
if( m_bOnEdge ) {m_bSuccess = FALSE; m_reason = 6; return;}
```

- 30

FAIL unconditionally largest object too small

```
m_pcntArea = m_ss[1].largestArea/ max(1.0, m_ss[0].largestArea);
```

```
if(m_pcntArea < m_pcntAreaMin ) { m_bSuccess = FALSE; m_reason = 5; return;}
```

FAIL unconditionally if too much clutter

```
if(m_ss[1].nObjects > m_nObjectsMax ) { m_bSuccess = FALSE; m_reason = 1;
return;}
```

5

FAIL unconditionally if shape skewed

```
m_dX1=m_ss[1].largestRect.right-m_ss[1].largestRect.left;
m_dY1=m_ss[1].largestRect.bottom-m_ss[1].largestRect.top;
m_formFactor = min(m_dX1, m_dY1) / max(1.0, max(m_dX1, m_dY1));
10 m_fillFactor = m_ss[1].largestArea / max(1.0, m_dX1 * m_dY1);
if(m_formFactor < m_formFactorMin ) { m_bSuccess = FALSE; m_reason = 2;
return;}
if(m_fillFactor < m_fillFactorMin ) { m_bSuccess = FALSE; m_reason = 3; return;}
```

15 SUCCEED if InvestigateNeighborhood comes up with a "good" {a,b,c} gradient

```
int xStart = ( m_ss[1].largestRect.right+m_ss[1].largestRect.left)/2;
int yStart = ( m_ss[1].largestRect.bottom+m_ss[1].largestRect.top )/2;
InvestigateNeighborhood(m_stamp, xStart, yStart, FALSE, &m_trCount1,
&m_trCount2,
20 &m_trCount3);
if( m_trCount1 >= m_trCount1Min1 )
if( m_trCount2 >= m_trCount2Min1 )
if( m_trCount3 >= m_trCount3Min1 )
{ m_bSuccess = TRUE; m_reason = 10; return; }
```

25

```
if( m_trCount1 >= m_trCount1Min2 )
if( m_trCount2 >= m_trCount2Min2 )
if( m_trCount3 >= m_trCount3Min2 )
{ m_bSuccess = TRUE; m_reason = 12; return; }
```

30

4. After each passed lesion

```
void CExpertSystem::UpdateDocStatistics()
```

```
switch( m_lesionType )
```

```
case IDOT:
```

```
(m_doc->m_fs.nDots)++;
```

```
(m_doc->m_fs.aveDotDef12) += (double)m_trCount1;
```

```
5 (m_doc->m_fs.aveDotDef23) += (double)m_trCount2;
```

```
(m_doc->m_fs.aveDotDef34) += (double)m_trCount3;
```

```
break;
```

```
Etc.
```

10 5. After each processed image (all lesions) write to mdb for postprocessing (Grading)

```
BOOL CDiagnosis::UpdateFieldMDB()
```

```
CFldResults pFR;
```

```
pFR.AddNew();
```

```
15 pFR.m_FieldCode = m_fCode;
```

```
pFR.m_sdGrayVal = m_doc->m_fs.sd;
```

```
pFR.m_meanGrayVal = m_doc->m_fs.mu;
```

```
pFR.m_minGrayVal = (BYTE)(m_doc->m_fs.minG);
```

```
pFR.m_maxGrayVal = (BYTE)(m_doc->m_fs.maxG);
```

```
20 pFR.m_nDots = m_doc->m_fs.nDots;
```

```
pFR.m_nBlots = m_doc->m_fs.nBlots;
```

```
pFR.m_nLipids = m_doc->m_fs.nLipids;
```

```
pFR.m_nCwools = m_doc->m_fs.nCwools;
```

```
pFR.m_aveDotDefinition =
```

```
25 100*(int)(m_doc->m_fs.aveDotDef12/max(1.0,m_doc->m_fs.nDots)+0.5)+
```

```
10*(int)(m_doc->m_fs.aveDotDef23/max(1.0,m_doc->m_fs.nDots)+0.5)+
```

```
(int)(m_doc->m_fs.aveDotDef34/max(1.0,m_doc->m_fs.nDots)+0.5);
```

```
pFR.m_aveBlotDefinition = 100*(int)(m_doc->m_fs.aveBlotDef12/max(1.0,m_doc-
```

```
30 >m_fs.nBlots)+0.5)+
```

```
10*(int)(m_doc->m_fs.aveBlotDef23/max(1.0,m_doc->m_fs.nBlots)+0.5)+
```

```
(int)(m_doc->m_fs.aveBlotDef34/max(1.0,m_doc->m_fs.nBlots)+0.5);
```



```

pFR.m_aveLipidDefinition = 100*(int)(m_doc-
>m_fs.aveLipidDef12/max(1.0,m_doc->m_fs.nLipids)+0.5)+
10*(int)(m_doc->m_fs.aveLipidDef23/max(1.0,m_doc->m_fs.nLipids)+0.5)+
(int)(m_doc->m_fs.aveLipidDef34/max(1.0,m_doc->m_fs.nLipids)+0.5);
5 pFR.m_aveCwoolDefinition = 100*(int)(m_doc-
>m_fs.aveCwoolDef12/max(1.0,m_doc->m_fs.nCWools)+0.5)+
10*(int)(m_doc->m_fs.aveCwoolDef23/max(1.0,m_doc->m_fs.nCWools)+0.5)+
(int)(m_doc->m_fs.aveCwoolDef34/max(1.0,m_doc->m_fs.nCWools)+0.5);
pFR.Update();
10
6. After mdb is updated for patient batch, all field results from the database (7/eye)
are unified to generate a grade between 1 & 3.

```

Quantify severity and use of Neural Net to generalize

- 15 Neural Net improves segmentation based on features extracted from the image
- * Histogram of number of lesions and size of lesions
 - * Types of lesions
- The best kind of matched filter is a classical neural network.
- 20 [a] For each lesion, pick an input layer sized to match the lesion size
- [b] Train on the areas in the target
- [c] Apply to the candidate regions identified during the false positive biased test.
- [d] There is no conceivable reason to identify more than one kind of lesion in a single pass, or to view the image as a single entity.
- 25 [e] The problem of classifying existence and number of lesions is fundamentally different from your dissertation because we don't need to recognize a "holistic" relationship or an "entity", just small objects with "clearly" defined properties in the bitmap. (The big objects are not relevant at this stage of the work, we can probably gloss over the features which require cross referencing amongst themselves).
- 30
- ENTITY (dynamic programming based segmentation) modified to accept a penalty function for the area of image to be parcelled); initializing penalty function (show with

colors in the annular region before starting PARCEL); defining prohibition zones (set penalty = 1)

Extraction of lesions given image and indication of field

5

Intelligent agent

Algorithm that learns params from images marked by a retina specialist. Associates params with image type.

- 10 This will help in having in one JAZZ, the following for 10 patients: the original, expert (dot/blot, lipid, striate, cotton wool spot), 4 comparisons and intermediate results

Implementation issues

- 15 Several polymorphic data structures representing regions and lines such that different representations enable us to write efficient algorithms for different kinds of processing. For example, a run-length encoded representation of a binary image allows fast determination of statistical properties like area, whereas an array representation allows one to access random cells faster for morphological processing.
- 20 Image Encoding Method for efficient lookup for recognition and interpretation

Physical implementation of the system is envisioned in the form of:

- I) Database of elementary images in SQL.
- II) Access through SQL to the Expert Validated Knowledge Base.
- 25 III) Database of scripts for finding each region in LISP: Rules for all Types of Knowledge.
- IV) Inference Engine in LISP: OTTO, OPS5, etc.
- V) Low level routines in C: Filters; ENTITY (dynamic programming based segmentation) modified to accept a penalty function for the area of image to be
- 30 parcelled); intializing penalty function (show with colors in the annular region before starting PARCEL); defining prohibition zones (set penalty = 1)

Comparison of sequential images

Comparison of sequential images to improve the risk prediction for pathologic changes occurring over time is provided by overlay and comparison of feature differences and by comparison in the database of the number and location of lesions (e.g. to detect new lesions or their migration toward anatomical features of interest, for instance toward the fovea).

While the present invention has been described above in terms of specific embodiments, it is to be understood that the invention is not limited to the disclosed embodiments. On the contrary, the present invention is intended to cover various modifications and equivalent structures included within the spirit and scope of the appended claims.

We claim as our invention:

1. A method for diagnosing diabetic retinopathy by computer implemented screening of retinal photographs comprising the steps of:

5 (A) receiving one or more images of a human retina and storing the retinal images in the memory of a digital computer;

(B) processing the stored original images to identify the presence of features indicative of diabetic retinopathy, the features of selected from the group of dot hemorrhages, microaneurysms, blot hemorrhages, striate hemorrhages, and lipid
10 exudates; and

(C) reporting the absence or presence and nature of the features indicative of diabetic retinopathy.

2. The method of claim 1 further comprising selecting those stored images for
15 processing that are deemed acceptable in terms of at least one image quality criterion selected from the group of discernible necessary elements, contrast, focus, alignment, and completeness.

3. The method of claim 1 wherein the received retinal images are captured
20 from a human retina by a retinal camera.

4. The method of claim 3 further comprising:

(D) assessing the stored retinal images in terms of at least one image quality criterion selected from the group of discernible necessary elements, contrast, focus,
25 alignment, and completeness; and

(E) if the assessment is unacceptable in any of the selected image quality criteria of step (D), then prompting an operator to capture further retinal image.

5. The method of claim 1 where in the retinal photographs are taken over one or more predetermined intervals and compared across one or more of the intervals to assess the development of the features indicative of diabetic retinopathy.

5 6. The method of claim 1 where in the original photographs are stored in the memory of the digital computer using an indexing convention that denotes the identity of the patient, the eye and field imaged, and the processing applied.

10 7. The method of claim 1 wherein one or more images are selected based on predetermined criteria and transmitted for examination by a human expert.

15 8. The method of claim 7 wherein the human expert is provided with both the transmitted images and the reported features indicative of retinopathy from step (C) of claim 1.

15 9. The method of claim 7 wherein the result of examination by the human expert is transmitted using the network and is stored in relation to the transmitted images.

20 10. A system for diagnosing diabetic retinopathy comprising a computer including a processor and memory wherein the processor is programmed to

(A) receive one or more images of a human retina and storing the retinal images in the computer memory;

25 (B) process the stored original images to identify the presence of features indicative of diabetic retinopathy, the features of selected from the group of dot hemorrhages, microaneurysms, blot hemorrhages, striate hemorrhages, and lipid exudates; and

(C) report the absence or presence and nature of the features indicative of diabetic retinopathy.

11. The system of claim 10 further comprising a retinal camera for capturing images from a human retina which is operatively linked to the computer.

5 12. The method of claim 10 further comprising at least one additional computer operatively linked to the computer using a network.

10 13. The method of claim 12 wherein one or more stored retinal images are selected based on predetermined criteria for transmission over the network for examination by a human expert.

15 14. The method of claim 13 wherein the human expert is provided with both the transmitted images and the reported features indicative of retinopathy from step (C) of claim 10.

15 15. The method of claim 13 wherein the result of examination by the human expert is transmitted using the network and is stored in relation to the transmitted images.




Figure 1. Histograms of two images of the fundus.




Figure 2 . (a) Initial image; (b) result of retina extraction



Figure 3. Result of processing the retina by coalescent filter [Bhasin and Meystel 94] to obtain the vessels

BEST AVAILABLE COPY




Figure 4. Problem with locating retina.




Figure 5. Insetting.




Figure 6(a) Subsampled original image




Figure 6(b) after subtraction of vessels from image

BEST AVAILABLE COPY




Figure 7. Results of finding the retina using geometry to improve the pickup

Overall Workflow




Fig. 8

Photographer Assist




Fig. 10

Serial Studies




Fig. 11